Government IT Business Continuity: To COOP or not to COOP?

By Elaine Price

You may remember the last scene in the movie, "Raiders of the Lost Ark." The ark of the covenant is wheeled, with little fanfare, into a vast government warehouse full of thousands upon thousands of boxes of various sizes. The warehouse stretches for what seems like miles; a cavern of sleeping data.

The federal government, and state and local governments across America, have more active and sleeping data than any enterprise in the history of mankind. The majority of this data deserves the title "mission-critical." Recognizing the importance of keeping this data protected and ensuring the continuous operation of systems in the event of disaster, the government makes heavy use of an acronym that has more currency today than ever: COOP, which stands for 'continuity of operations planning.'

The Federal Emergency Management Agency (FEMA) first formally recognized COOP with Federal Preparedness Circular 65, published in July of 1999. The purpose of the circular was to provide "guidance to federal executive branch departments and agencies for use in developing viable and executable contingency plans for the continuity of operations. COOP planning facilitates the performance of department/agency essential functions during any emergency or situation that may disrupt normal operations."

In April 2003, R. Eric Petersen wrote an important follow-up report entitled "COOP in the Executive Branch: Background and Issues for Congress." The report reads: "In the wake of the September 2001 terrorist attacks … federal policymakers have given renewed attention to continuity of operations (COOP) issues. COOP planning is a segment of federal government contingency planning that refers to the internal effort of an organization … to assure that the capability exists to continue essential operations in the aftermath of a comprehensive array of potential operational interruptions. Governmentwide, COOP planning is critical because much of the recovery from an incident … presumes the existence of an ongoing, functional government to fund, support, and oversee actions taken. COOP planning can be viewed as a continuation of basic emergency preparedness planning, and a bridge between that planning and efforts to maintain continuity of government in the event of a significant disruption to government activity or institutions. Because the number and types of potential interruptions are unknown, effective COOP planning must provide, in advance of an incident, a variety of means to assure contingent operations."

Peterson's report is an indication of prevailing COOP strategies that increasingly focus on keeping government IT systems running with 100 percent uptime and ensuring that data is well-protected. But are these words being projected into actions, and if so, how? The answer is yes. Government agencies, along with their civilian counterparts, are satisfying COOP requirements with four broad approaches that are very much "COTS," or commercial-off-the-shelf solutions:

- Better backup
- Increased replication

- The advent of electronic records management
- Intelligent archiving

Better Backup
Traditional backup software has been around for almost 30 years. Typically, it is run nightly to create a snapshot of changes that occurred during the day, and is amplified with a weekly full backup. The resulting tape cartridges are then either sent offsite to a secure location, or stored in a fireproof safe on the premises.

These traditional backup solutions have been generic in regard to the over-arching application — the same for Peoplesoft, Lotus Notes, and Documentum applications.

However, generic backup solutions are no longer the only piece of the backup puzzle required to ensure that mission-critical data is safe. A new breed of "application-aware" backup software vendors are delivering powerful, advanced functionality by drilling into the application programming interface (API) of the over-arching application, and intimately understanding its data schema. These solutions run "hot," such that the application does not need to be brought down. Backups run incrementally, at user-defined intervals as frequently as every hour or more, dramatically narrowing the worst-case data loss window.

Perhaps best of all, application-aware backup solutions allow for granular, object-level or message-level (in the case of e-mail) restores. Instead of having to restore a volume of data to recover the small element that was lost, the actual element itself can rapidly be recovered.

Further, disk is replacing tape as the backup target of choice. Disk-to-disk backup, commonly referred to as D2D, is far faster to restore from than tape, and tape cartridges suffer data loss as they age. There is still an important role for tape given its portability, but increasingly organizations are keeping a minimum of 6 months of incremental backups and a few recent full backups on local RAID disk subsystems.

This is an area to watch in the next nine to 12 months, as competition between existing tape companies and emerging disk vendors will soon grow fierce. The disk vendors are using low-cost ATA disk drives to hit disk subsystem prices similar to tape. Disk vendors are now working to make the disk drives portable, such that a drive can be pulled for offsite storage.

Increased Replication
Despite better backup solutions, there is a problem with "just backing up." Although data and system backup is a key piece of the overall COOP puzzle, government organizations that don't consider replication will be ill prepared to maintain continuity in many disaster scenarios whereby the computer system and application can be wiped out. In this scenario, tapes stored offsite are of little good since getting them to the proper location can take days, depending on the nature of the disaster. In addition, any data that was created or generated between the time of the last backup tape being shipped out and the disaster will have been lost. In some cases that window may be greater than a week. Thus, many organizations employ either synchronous or asynchronous replication, which means making a copy of the data with some level of simultaneity to a mirror system in another (usually somewhat distant) geographic location.

Given that servers and storage cost less by the year, replication is becoming more affordable. But again, it is application-awareness that is making replication more palatable in terms of cost and ease of operation. Increasingly, enterprise software programs have replication built-in, so it can be carried out at the application level, rather than file level. If this functionality is not built-in, it is available from third-party software vendors.

A December 1, 2003 article in The Washington Post points out the benefits of replication, as follows: "Virginia's top emergency official said Monday that a computer designed to track requests for help from local governments failed repeatedly during Hurricane Isabel, delaying the distribution of ice, water, generators and other assistance. Michael Cline, the director of the state's Department of Emergency Management, told lawmakers at a hearing that the state's 'Action Tracking' computer system did not have a battery backup. When the power went out, requests for help were lost."

The Advent of Electronic Records Management (ERM)
Government entities have long realized the value of their data and the importance of uptime. But given their gargantuan data volumes, storing, protecting and managing the millions of records and database rows has been a significant challenge.

One part of the problem is that not all data is equal — some must be kept forever, some only for a year; some data needs maximum protection while other data is not as essential. There is a clear lifecycle of data from birth to death, but until recently there has been a shortage of software tools to manage this lifecycle.

Now ERM and ILM (information lifecycle management) software is proliferating, such that data can easily be tagged in terms of how long to retain it, among other factors. Applications are building ERM /ILM capabilities in natively, or obtaining this functionality from formidable third-party applications. Using ERM/ILM, the classification and categorization of volumes of data can be done automatically based on a set of pre-defined rules, making if far less unwieldy.

Intelligent Archiving
Even with ERM and ILM, the overwhelming size of government data results in enormous storage costs, and this has always been part of the problem with keeping applications up and protected. Many government installations, such as in the arena of satellite imaging, have applications with literally hundreds of terabytes — even petabytes — of data. Keeping this data online and accessible has been almost impossible. Relying on disk alone is far too expensive, as are the various hierarchical storage management (HSM) strategies relied upon to migrate infrequently accessed data to lower-cost tape or optical media. This was fine in terms of lowering the cost of storage, but the robotic arms associated with tape and optical jukeboxes proved to be a potentially disastrous single-point-of-failure, and the jukeboxes could not sustain even a moderate level of concurrent requests without response times slowing to a snail's pace.

Add to all of this the fact that some data had to be archived in an unalterable format, which requires the use of some type of WORM (write once-read many) media.

However, today's disk vendors are using ATA drives to create 10TB storage systems that cost about the same as a tape or optical library. These systems are being used

for backup and for archiving. Infrequently accessed data is thus migrated from expensive, "fast" disk to this lower-cost, "slow" disk. There are no robotics, and these systems can sustain massive concurrent usage without response-time degradation. Further, these systems can be WORM — that is, the drives can be rendered unalterable or "nonrewritable" through software inside the disk subsystem.

There's no better time to address COOP issues in government than today. A variety of vendors are delivering commercial off-the-shelf hardware and software solutions for better backup, replication, archiving and classifying of data. Government agencies can now find a way to keep their applications continually operating in a world where the risks and costs of data loss have never been higher.

About the author: Elaine Price is the president and CEO of CYA Technologies.